Trustworthy Networking Program

Architectural Considerations for Mapping Distribution Protocols

Kotikapaludi Sriram, Young-Tak Kim, and Doug Montgomery

National Institute of Standards and Technology

Contact: ksriram@nist.gov

Outline of the Talk

- **Problem statement**

- Management of mapping distribution of subprefixes spread across multiple ETRs

- Management of mapping distribution for scenarios with a hierarchy of ETRs and multi-homing

- Conclusions / Future Work

Problem Statement

- Aggregation of EID address space can be highly beneficial for the mapping distribution system

- But exceptions/violations will occur

- Some subprefixes of an aggregatable prefix would be often located in other parts of the world

- What can be done to handle these exceptions in the mapping distribution protocols?

- Other types of exceptions: Scenarios with Multi-homing and hierarchy of ETRs

- Goal: Better efficiency/performance of mapping distribution system
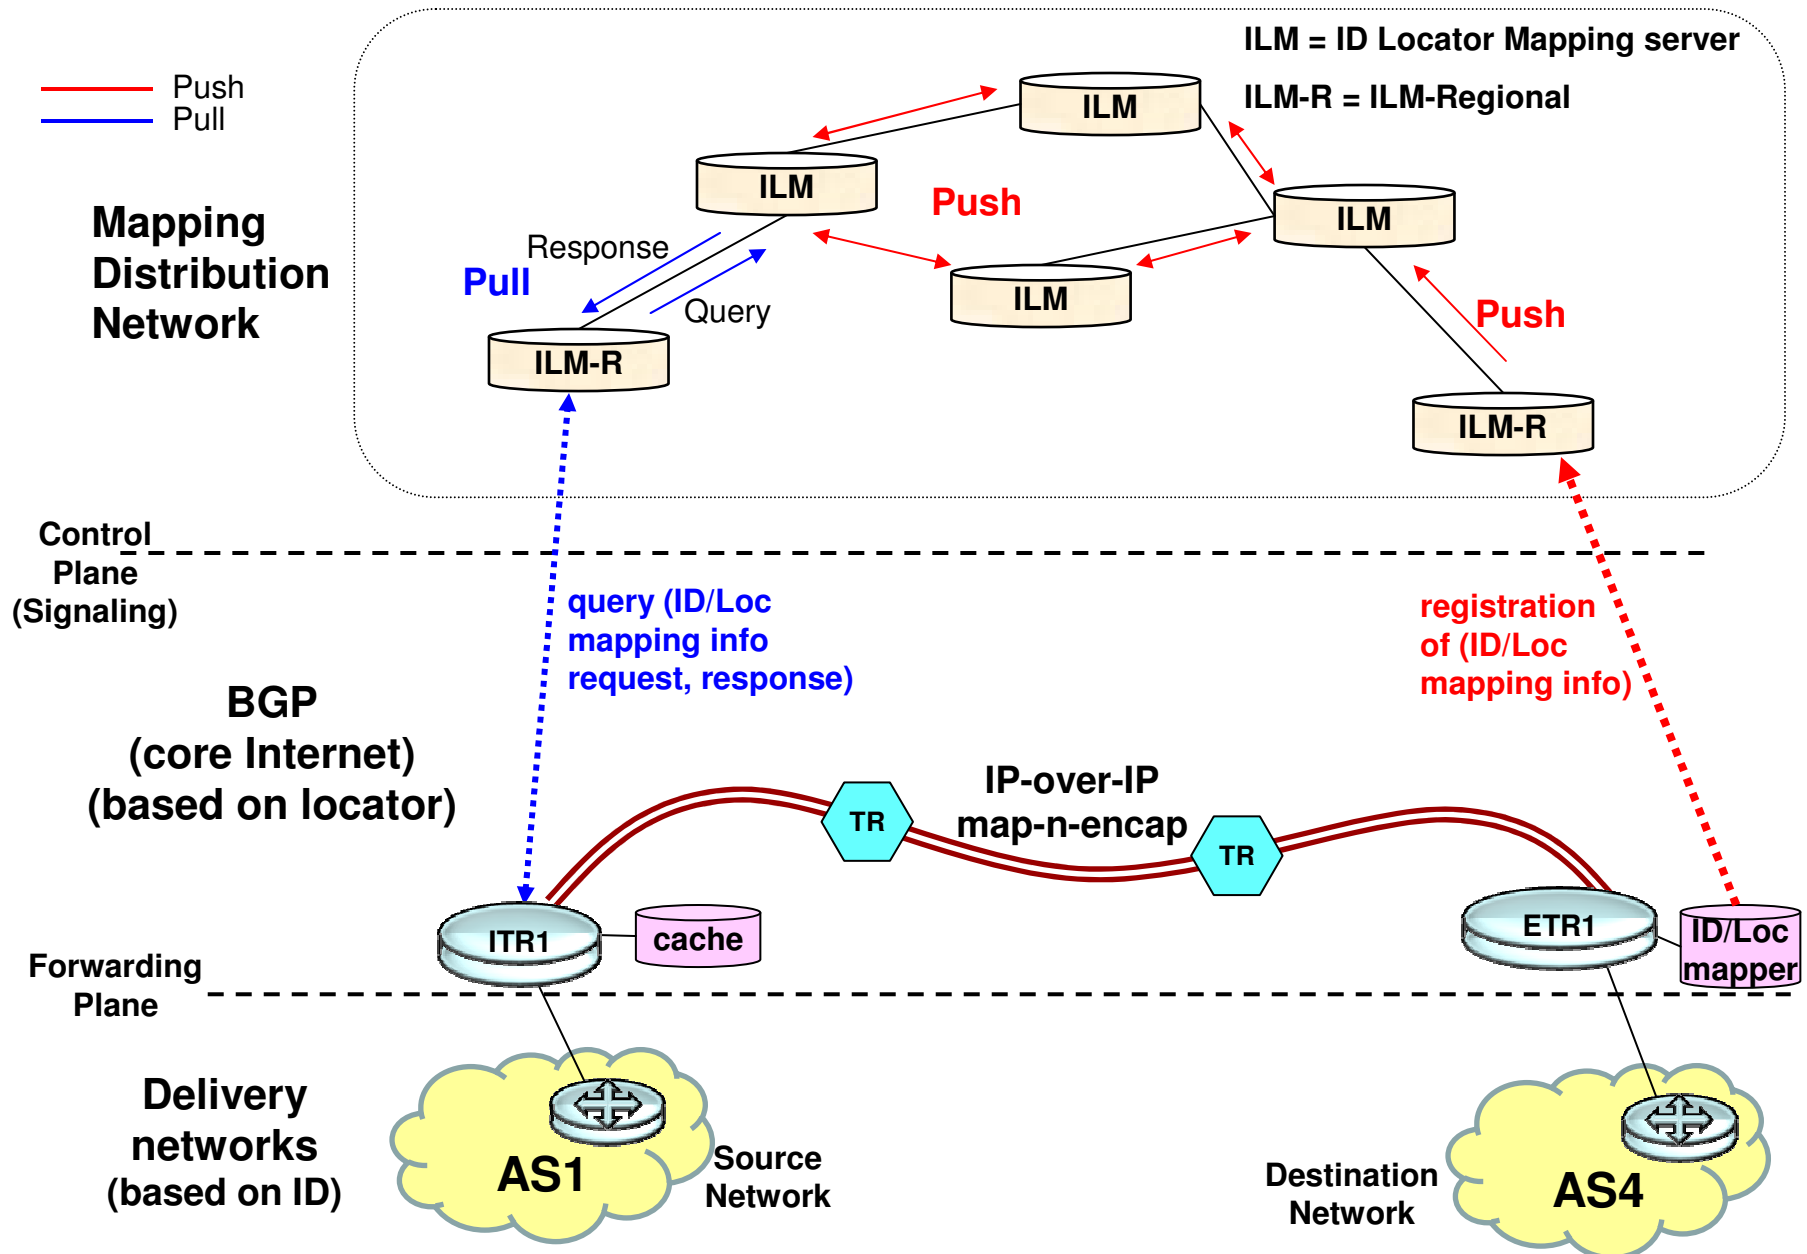
Quoting Brian Dickson (July 9, 2008)*:

- Basically, scaling requires "layering", e.g. aggregation, or indirection, or other similar techniques.

- But, a dynamic location in the food chain, automatically means that we need the **ability** to routinely handle "layer violations".

- I think drilling down on this, providing means to incorporate such violations **into** the addressing /identifier /locator/ mapping system, is the way forward.

http://psg.com/lists/rrg/2008/msg01773.html

* I noticed Brian's email well after I submitted this contribution to RRG list
http://www.antd.nist.gov/~ksriram/NGRA_map_mgmt.pdf

Overview of an Example Architecture

Subprefixes Spread Across Multiple ETRs

ILM = ID Locator Mapping server

ILM-R = ILM-Regional

a_8/24

a_{15}/24

ETR3

ETR2

ILM

ILM

ILM

ILM

ILM-R

ILM-R

ITR1

ETR1

Parent ETR for a/20

[a_1/24, … , a_7/24, a_9/24, … a_{14}/24, a_{16}/24]
(Note: Most of a/20 is here; except subprefixes
a_8/24 and a_{15}/24 which are elsewhere)
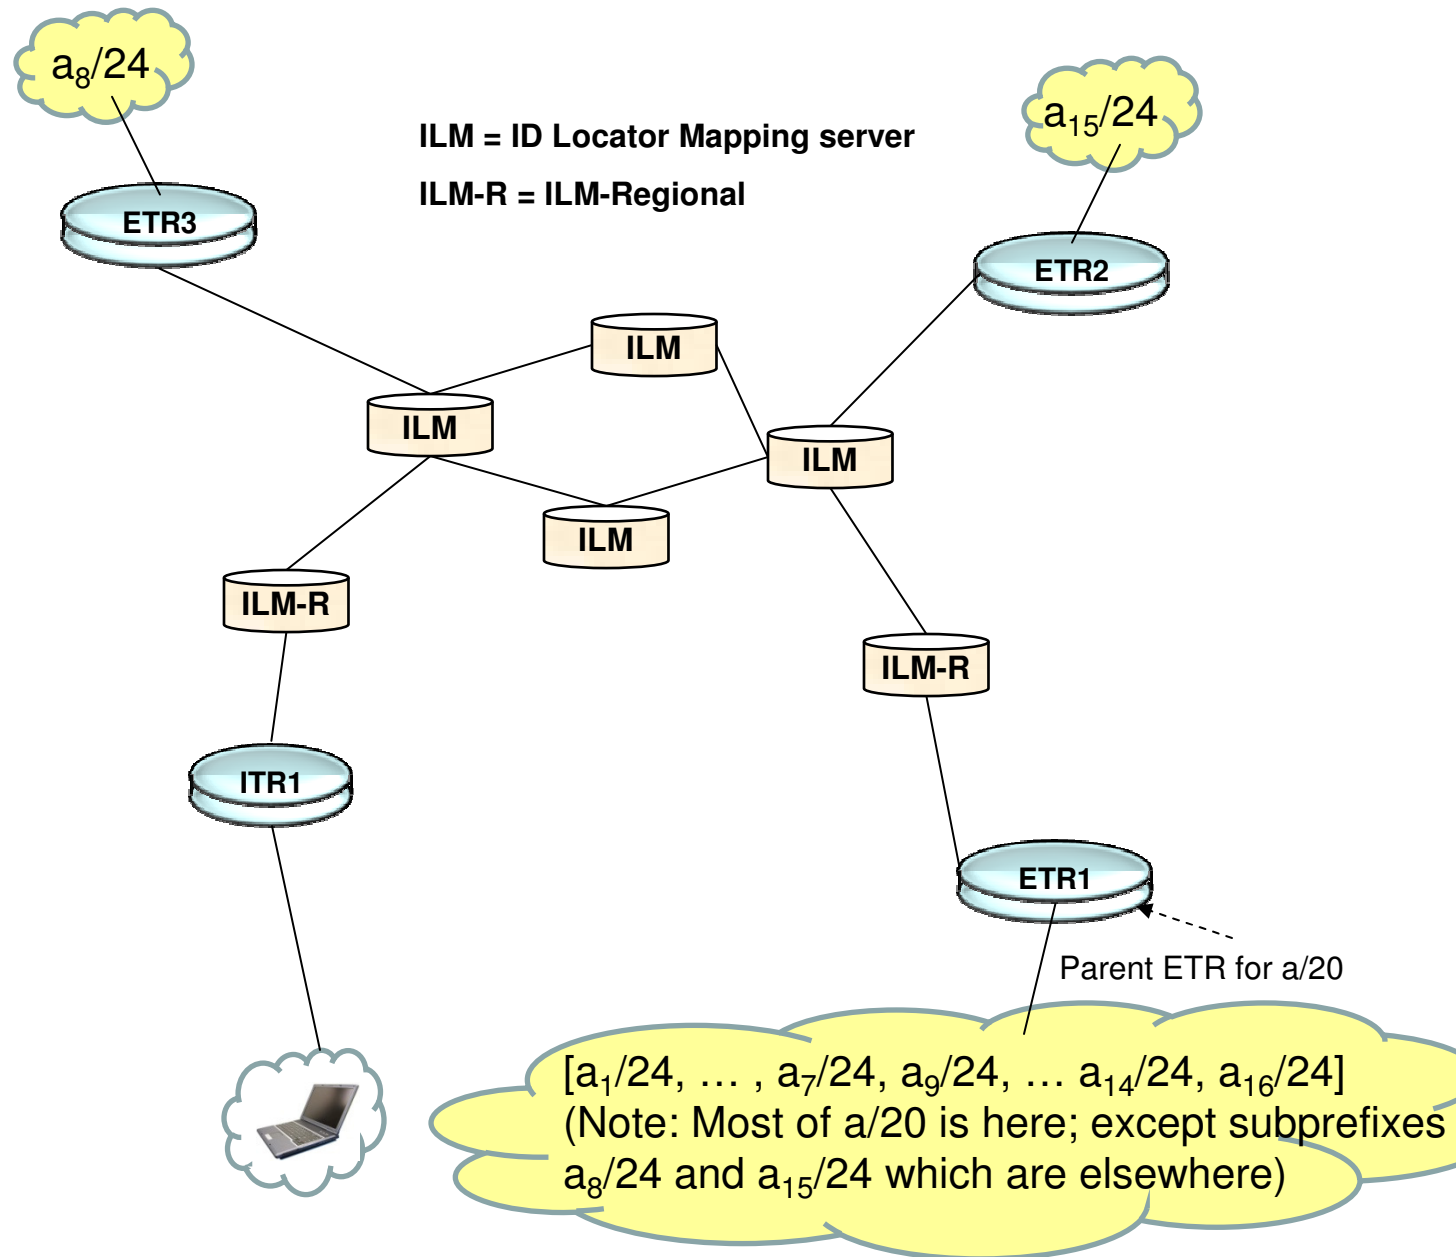
Possible Solution/Limitation in LISP-ALT
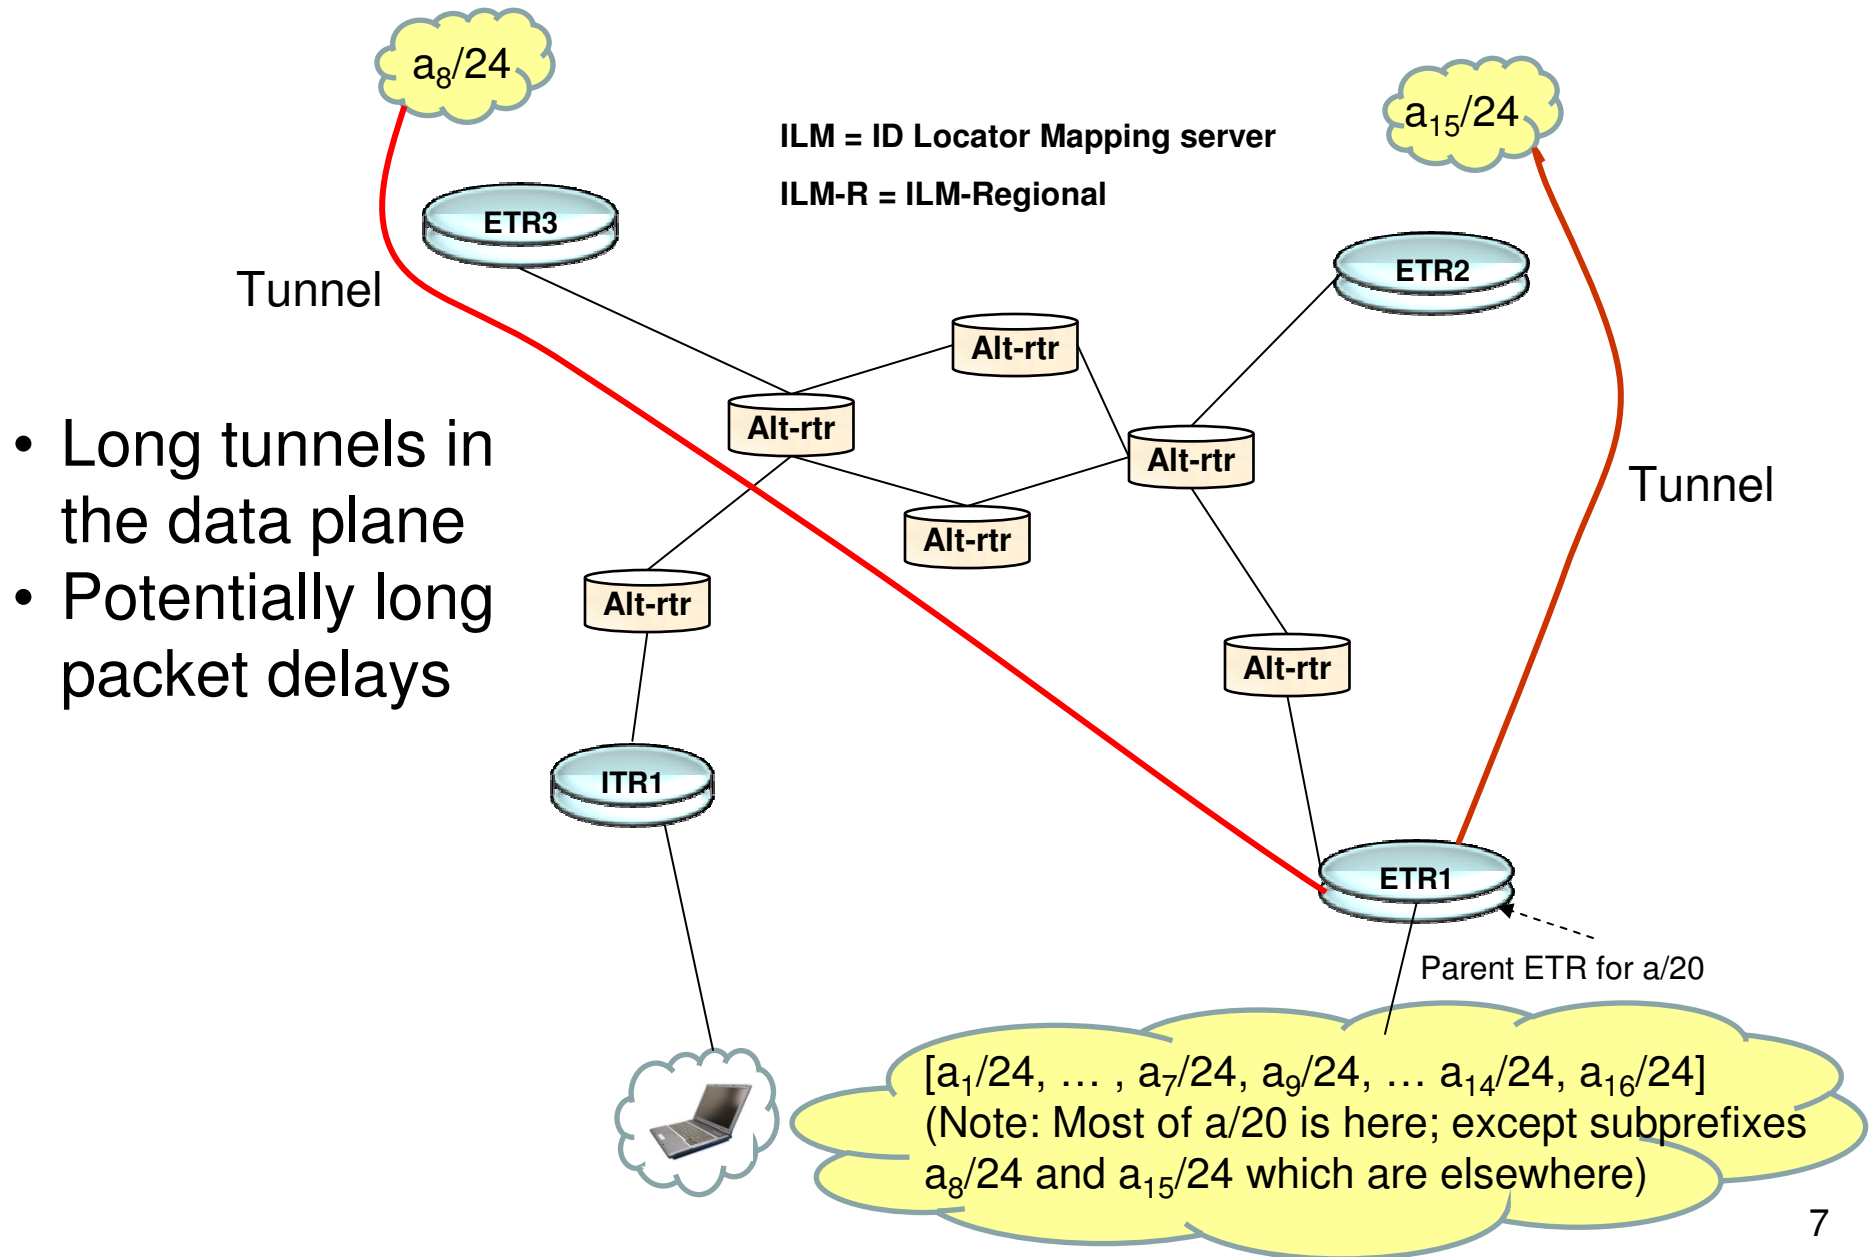
a_8/24

a_{15}/24

ILM = ID Locator Mapping server

ILM-R = ILM-Regional

ETR3

ETR2

Tunnel

Alt-rtr

Alt-rtr

Alt-rtr

- Long tunnels in the data plane
- Potentially long packet delays

Alt-rtr

Alt-rtr

Tunnel

Alt-rtr

ITR1

ETR1

Parent ETR for a/20

[a_1/24, ... , a_7/24, a_9/24, ... a_{14}/24, a_{16}/24]
(Note: Most of a/20 is here; except subprefixes a_8/24 and a_{15}/24 which are elsewhere)

7

Outline of the Talk

- **Problem statement**

- **Management of mapping distribution of subprefixes spread across multiple ETRs**

- Management of mapping distribution for scenarios with a hierarchy of ETRs and multi-homing

- Conclusions / Future Work

Subprefixes Spread Across Multiple ETRs

Push

$a_8/24$

$a_{15}/24$

ILM = ID Locator Mapping server

ILM-R = ILM-Regional

ETR3

ETR2

ILM

ILM

ILM

ILM

Map: ($a_8/24$, ETR3)

Map: ($a_{15}/24$, ETR2)

ILM-R

ILM-R

ITR1

Map: (a/20, ETR1,
Exceptions: $a_8/24$, $a_{15}/24$)

ETR1

Parent ETR for a/20

[$a_1/24$, … , $a_7/24$, $a_9/24$, … $a_{14}/24$, $a_{16}/24$]
(Note: Most of a/20 is here; except subprefixes
$a_8/24$ and $a_{15}/24$ which are elsewhere)

9

Subprefixes Spread Across Multiple ETRs

$a_8/24$

Pull

ILM = ID Locator Mapping server

ILM-R = ILM-Regional

$a_{15}/24$

ETR3

ETR2

ILM

ILM

ILM

ILM

ILM

Query: a_7^1

ILM-R

ILM-R

Query: a_7^1

ITR1

ETR1

Packet destined to an address a_7^1 in $a_7/24$

Parent ETR for a/20

$[a_1/24, \ldots , a_7/24, a_9/24, \ldots a_{14}/24, a_{16}/24]$
(Note: Most of a/20 is here; except subprefixes $a_8/24$ and $a_{15}/24$ which are elsewhere)
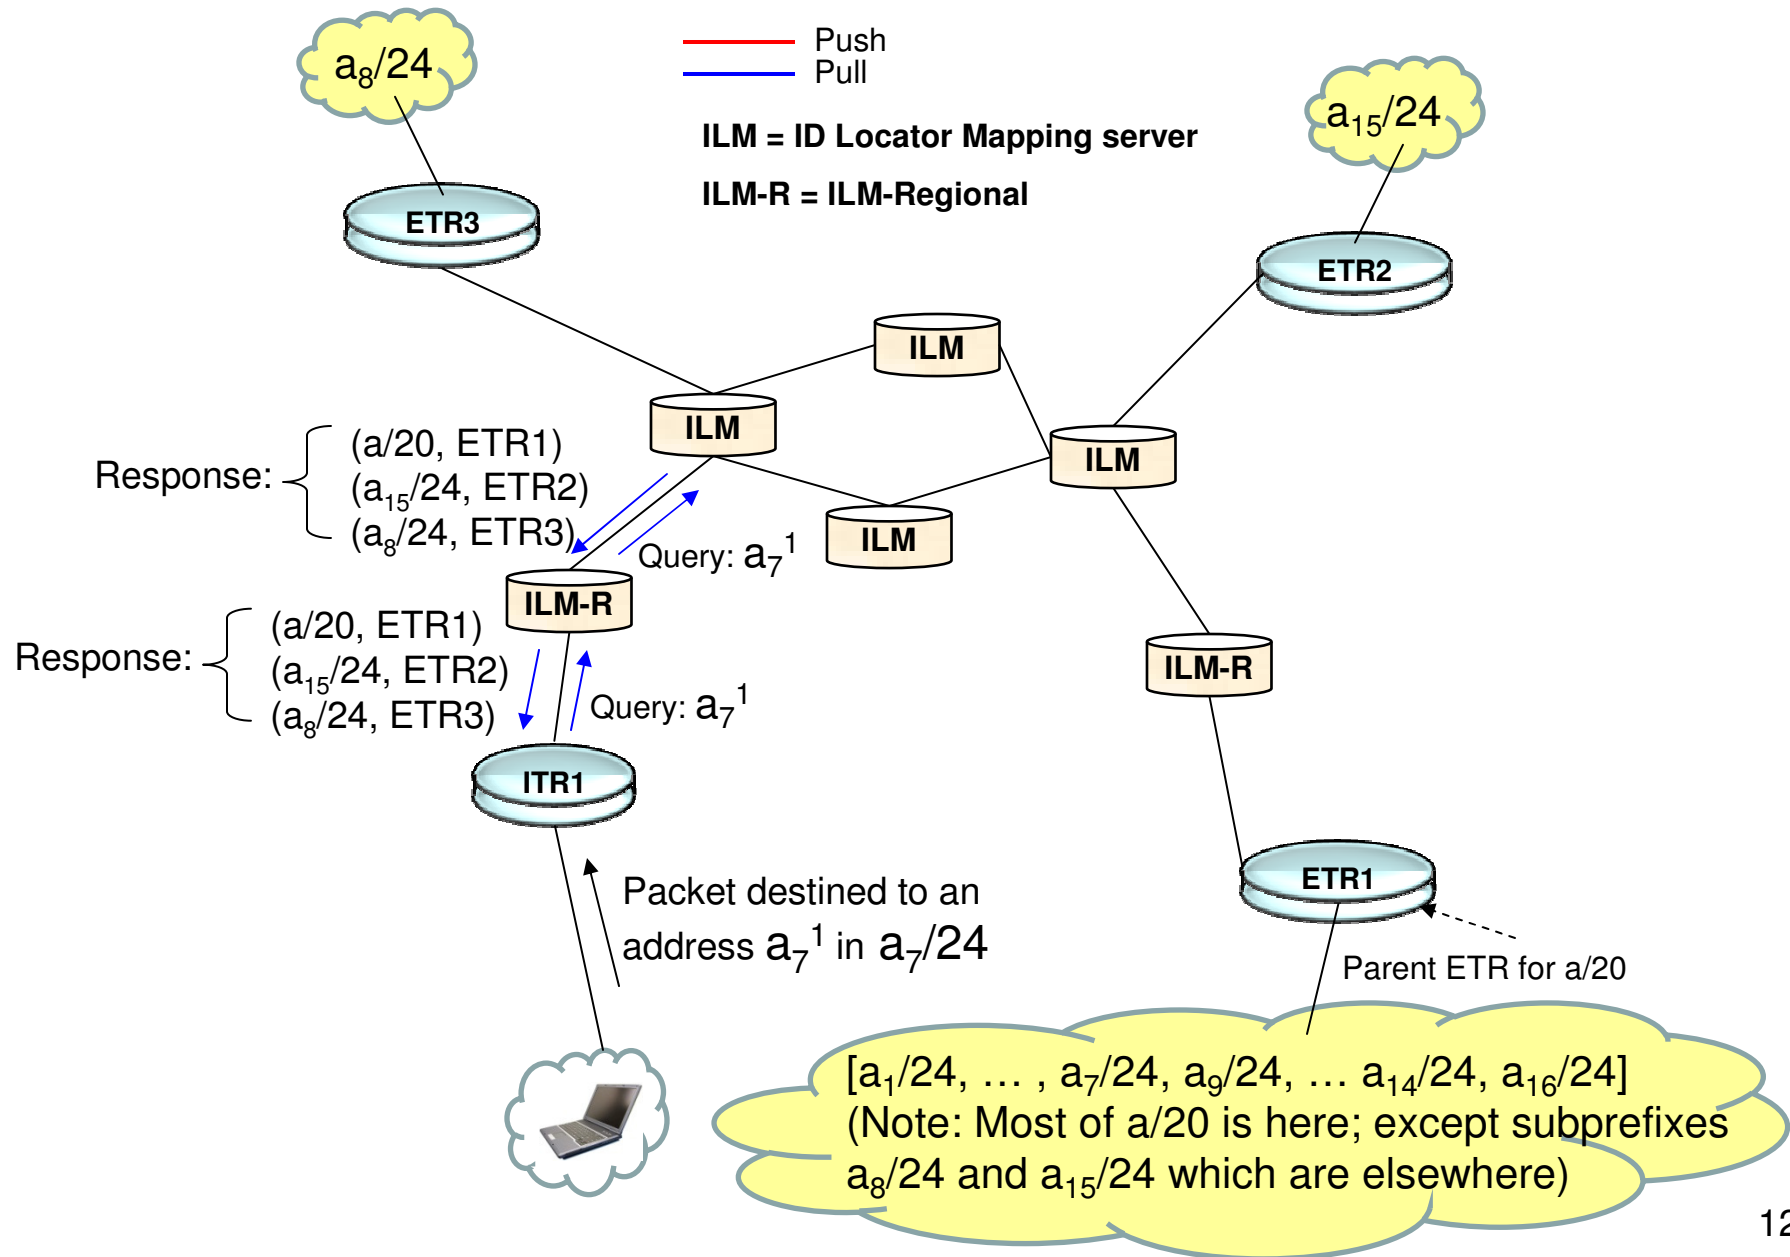
Alternative Ways of
ILM-R to ITR Mapping Distribution

Approach 1: ILM-R provides the complete mapping information for $a/20$ to ITR1 including all the maps for the relevant exception subprefixes.

Approach 2:

Approach 3:

Mapping Distribution: Approach 1

Push (red line)
Pull (blue line)

ILM = ID Locator Mapping server

ILM-R = ILM-Regional

$a_8/24$

ETR3

$a_{15}/24$

ETR2

ILM

ILM

ILM

Response:
$(a/20, ETR1)$
$(a_{15}/24, ETR2)$
$(a_8/24, ETR3)$

ILM

Query: a_7^1

ILM-R

Response:
$(a/20, ETR1)$
$(a_{15}/24, ETR2)$
$(a_8/24, ETR3)$

ILM-R

Query: a_7^1

ITR1

ETR1

Packet destined to an address a_7^1 in $a_7/24$

Parent ETR for a/20

$[a_1/24, \ldots, a_7/24, a_9/24, \ldots a_{14}/24, a_{16}/24]$
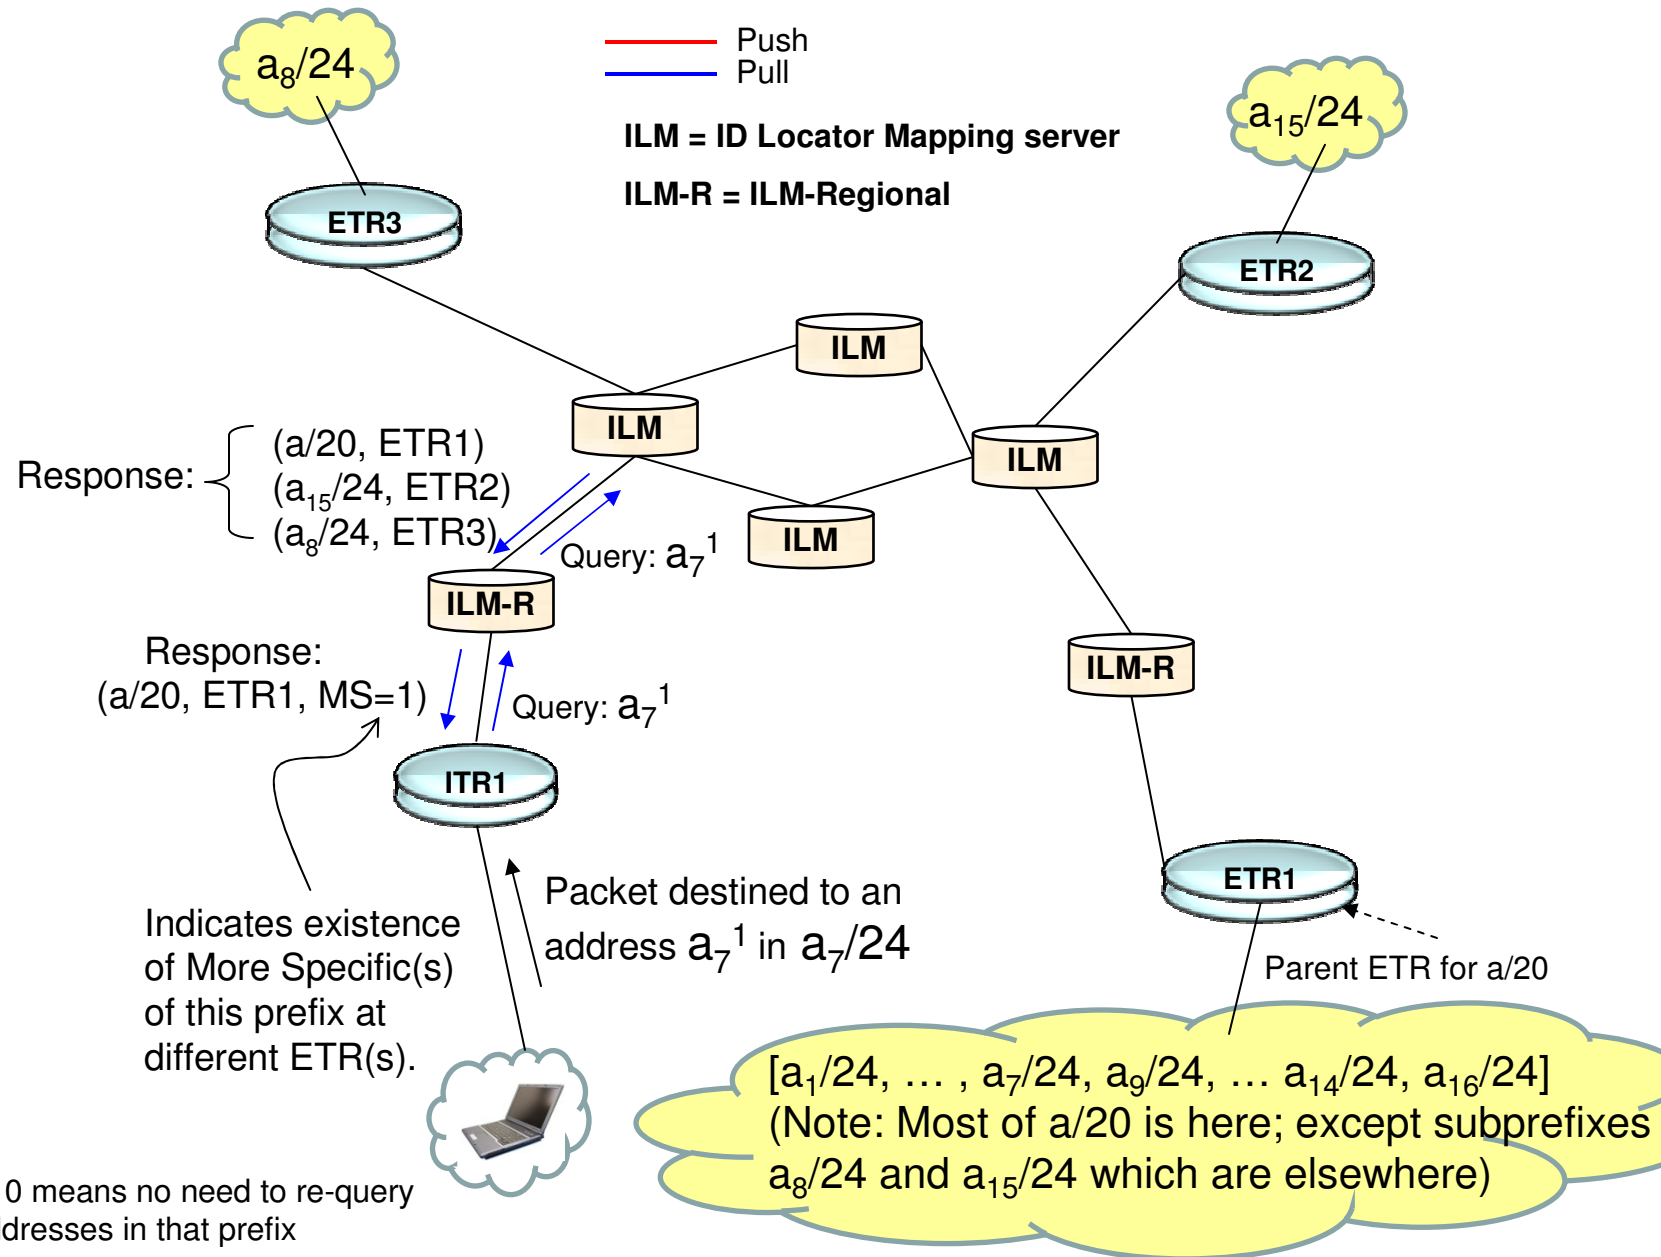(Note: Most of a/20 is here; except subprefixes $a_8/24$ and $a_{15}/24$ which are elsewhere)

12

Mapping Distribution: Approach 2

Approach 1: ILM-R provides the complete mapping information for *a*/20 to ITR1 including all the maps for the relevant exception subprefixes.
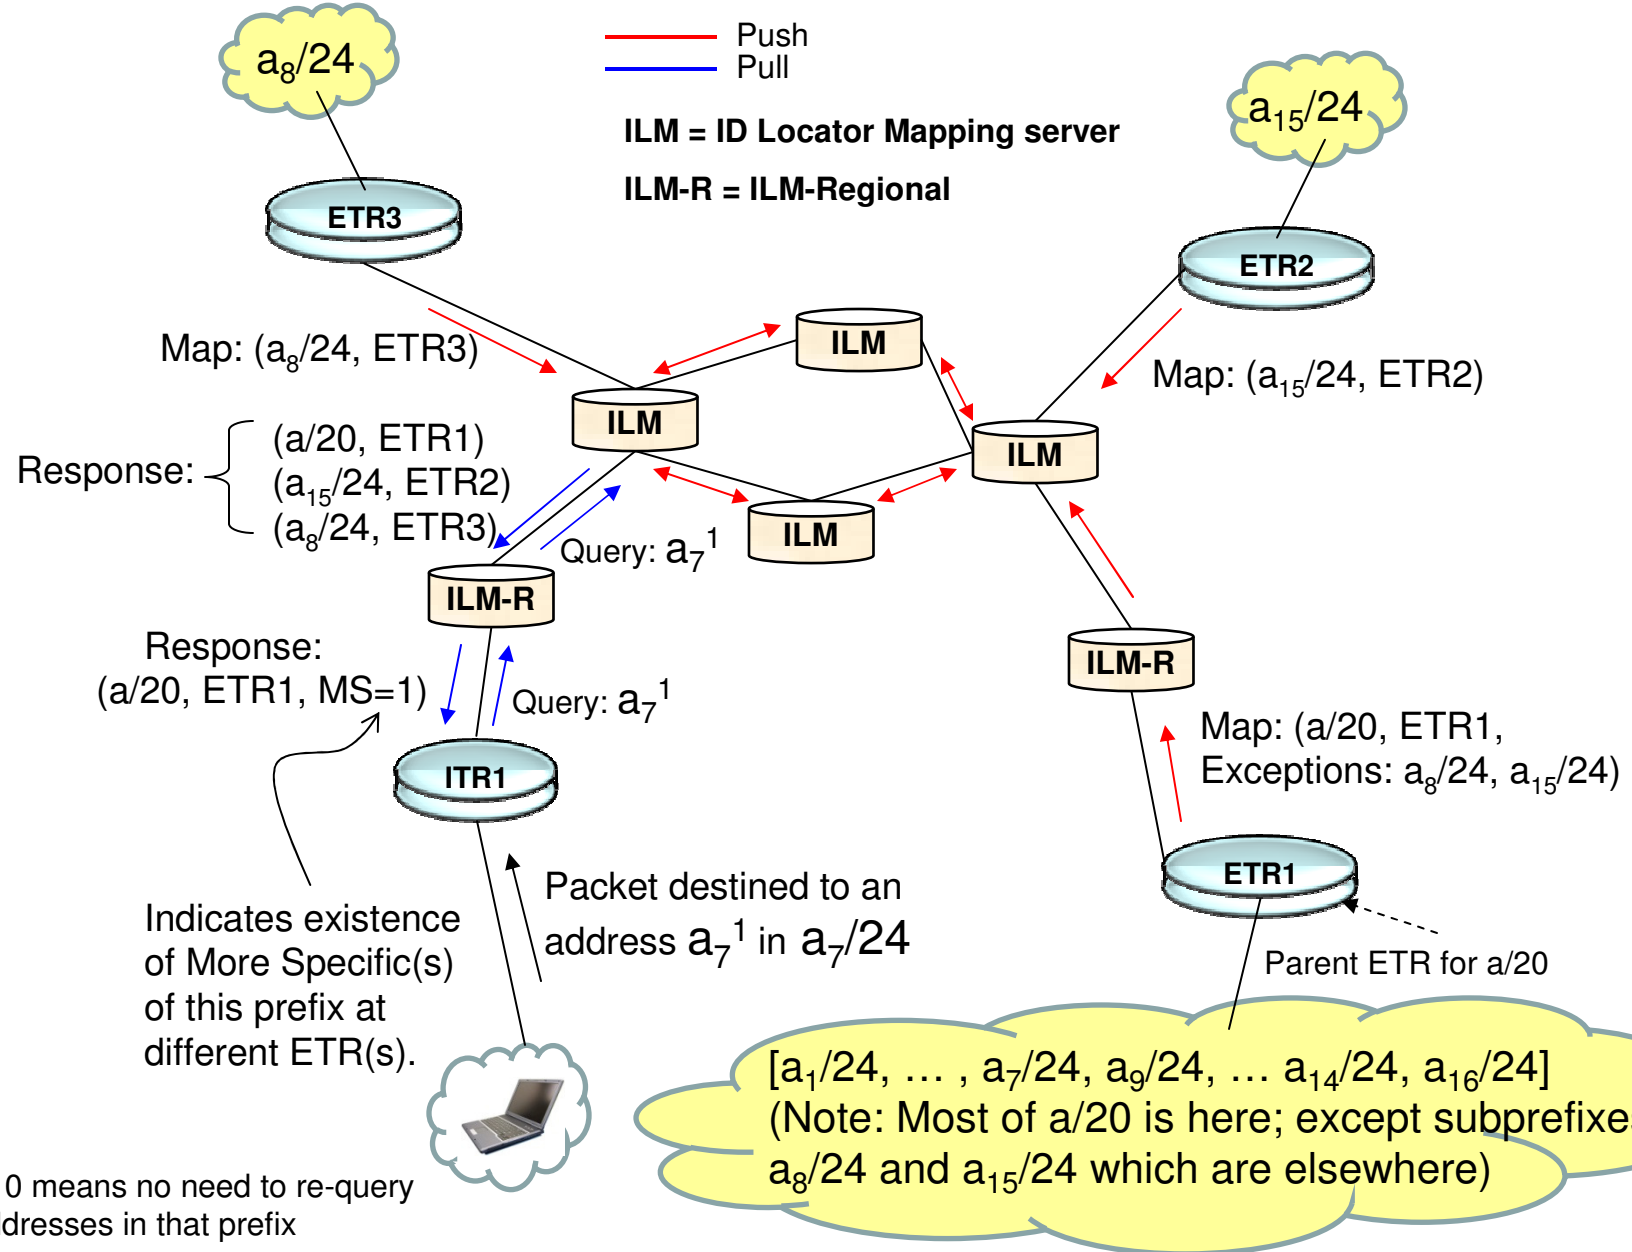
Approach 2: ILM-R provides only the directly relevant map to ITR1 which in this case is *(a/*20, ETR1).

Approach 3:

Mapping Distribution: Approach 2

$a_8/24$

Push
Pull

ILM = ID Locator Mapping server

ILM-R = ILM-Regional

$a_{15}/24$

ETR3

ETR2

ILM

ILM

ILM

ILM

Response:
- (a/20, ETR1)
- $(a_{15}/24, \text{ETR2})$
- $(a_8/24, \text{ETR3})$

Query: a_7^1

ILM-R

ILM-R

Response:
(a/20, ETR1, MS=1)

Query: a_7^1

ITR1

ETR1

Indicates existence
of More Specific(s)
of this prefix at
different ETR(s).

Packet destined to an
address a_7^1 in $a_7/24$

Parent ETR for a/20

$[a_1/24, \ldots, a_7/24, a_9/24, \ldots a_{14}/24, a_{16}/24]$
(Note: Most of a/20 is here; except subprefixes
$a_8/24$ and $a_{15}/24$ which are elsewhere)

MS = 0 means no need to re-query
for addresses in that prefix

14

Full Diagram with All Flows: Approach 2

$a_8/24$

— Push
— Pull

ILM = ID Locator Mapping server

ILM-R = ILM-Regional

$a_{15}/24$

ETR3

ETR2

Map: ($a_8/24$, ETR3)

ILM

Map: ($a_{15}/24$, ETR2)

ILM

Response:
  (a/20, ETR1)
  ($a_{15}/24$, ETR2)
  ($a_8/24$, ETR3)

ILM

ILM

Query: a_7^1

ILM

ILM-R

Response:
(a/20, ETR1, MS=1)

Query: a_7^1

ILM-R

ITR1

Map: (a/20, ETR1,
Exceptions: $a_8/24$, $a_{15}/24$)

Indicates existence
of More Specific(s)
of this prefix at
different ETR(s).

Packet destined to an
address a_7^1 in $a_7/24$

ETR1

Parent ETR for a/20

[$a_1/24$, … , $a_7/24$, $a_9/24$, … $a_{14}/24$, $a_{16}/24$]
(Note: Most of a/20 is here; except subprefixes
$a_8/24$ and $a_{15}/24$ which are elsewhere)

MS = 0 means no need to re-query
for addresses in that prefix

15

Pros/Cons Approach 1

ILM-R provides the complete mapping information for *a*/20 to ITR1 including all the maps for the relevant exception subprefixes.

➢ **Advantages: ITR1 has all the complete Mapping information for all components of *a*/20 and need not generate repeated queries for subsequent first packets; Better response time for the first packets**

➢ **Disadvantage: If the exception subnets are significant in number, then both ITR1 and ILM-R can face processing overload; Also, results in inefficient usage of cache/memory at ITR1**

Pros/Cons Approach 2

ILM-R provides only the directly relevant map to ITR1 which in this case is (a/20, ETR1).

➢ **Advantages: This will help avoid resource exhaustion at the ITRs and also possibly at ILM-R**

➢ **Disadvantages:**

  ▪ **ITR1 must be informed that there are potentially more specifics (via use of More Specific (MS) field in the Map)**

  ▪ **ITR1 needs to re-enquire for each first packet**

  ▪ **Slightly larger first packet delays overall**

Mapping Distribution: Approach 3

Approach 1: ILM-R provides the complete mapping information for *a*/20 to ITR1 including all the maps for the relevant exception subprefixes.

Approach 2: ILM-R provides only the directly relevant map to ITR1 which in this case is *(a/*20, ETR1).

Approach 3: The mapping information transaction between ILM-R and ITR1 can dynamically use approach 1 or approach 2 above depending on the context (further explanation of this is provided below).

Benefits of Approach 3

The mapping information transaction between ILM-R and ITR1 can dynamically use approach 1 or approach 2 above depending on the context.

➢ Here a parameter can be potentially defined such as the maximum number of maps (map for the parent prefix plus all the maps for the relevant exceptions subprefixes) that would be involved in a mapping transaction

➢ If this parameter is below a threshold, then the first approach would be used, else, the second approach would be used

➢ This parameter can also be tuned administratively or dynamically (e.g., dependent on load at ILM-R)

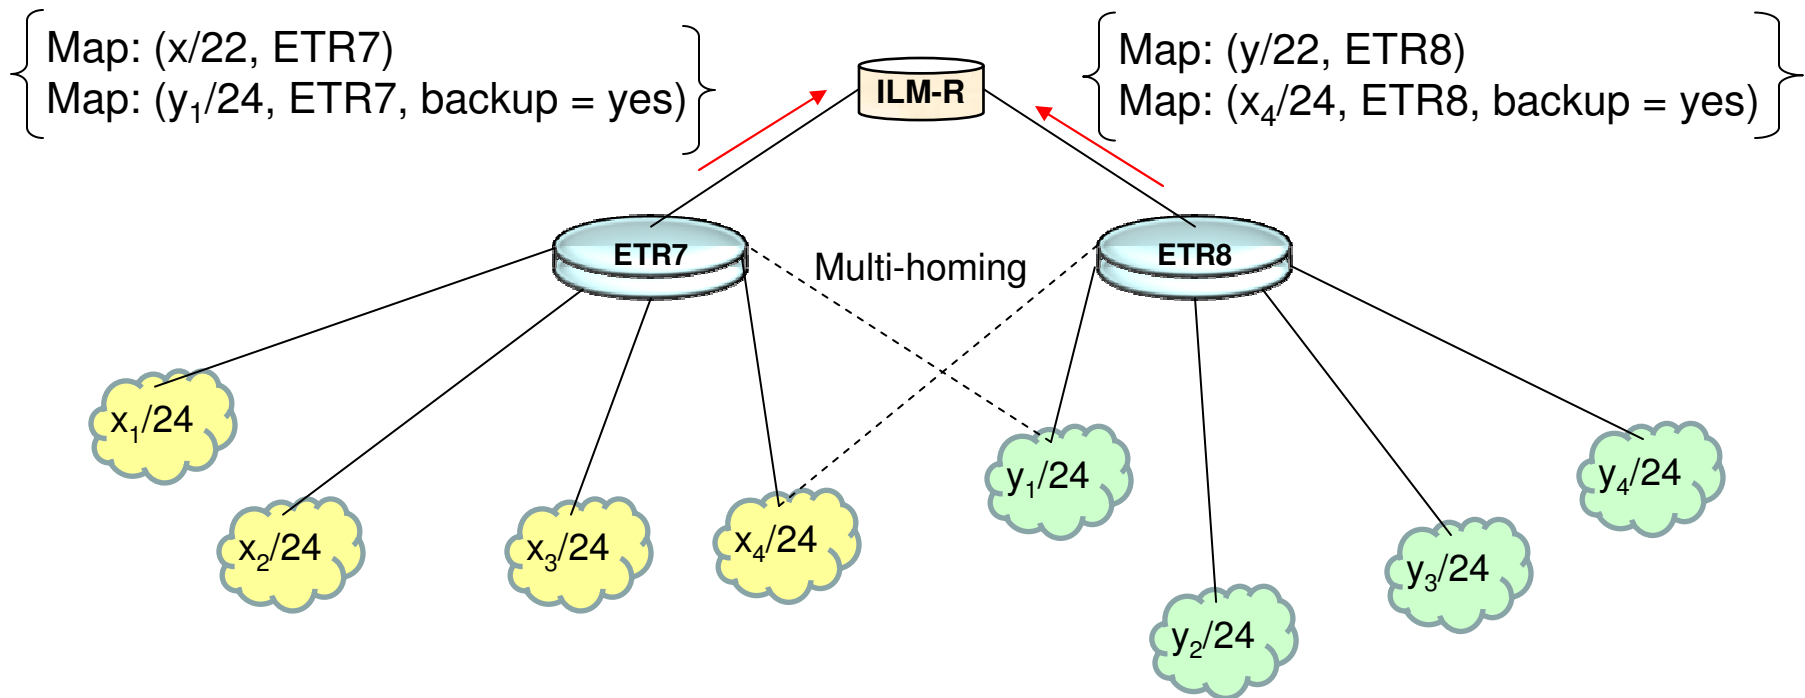➢ Seeking to combine of benefits of approaches 1 & 2

Advantage of Approach 3
(Conceptual Plot)
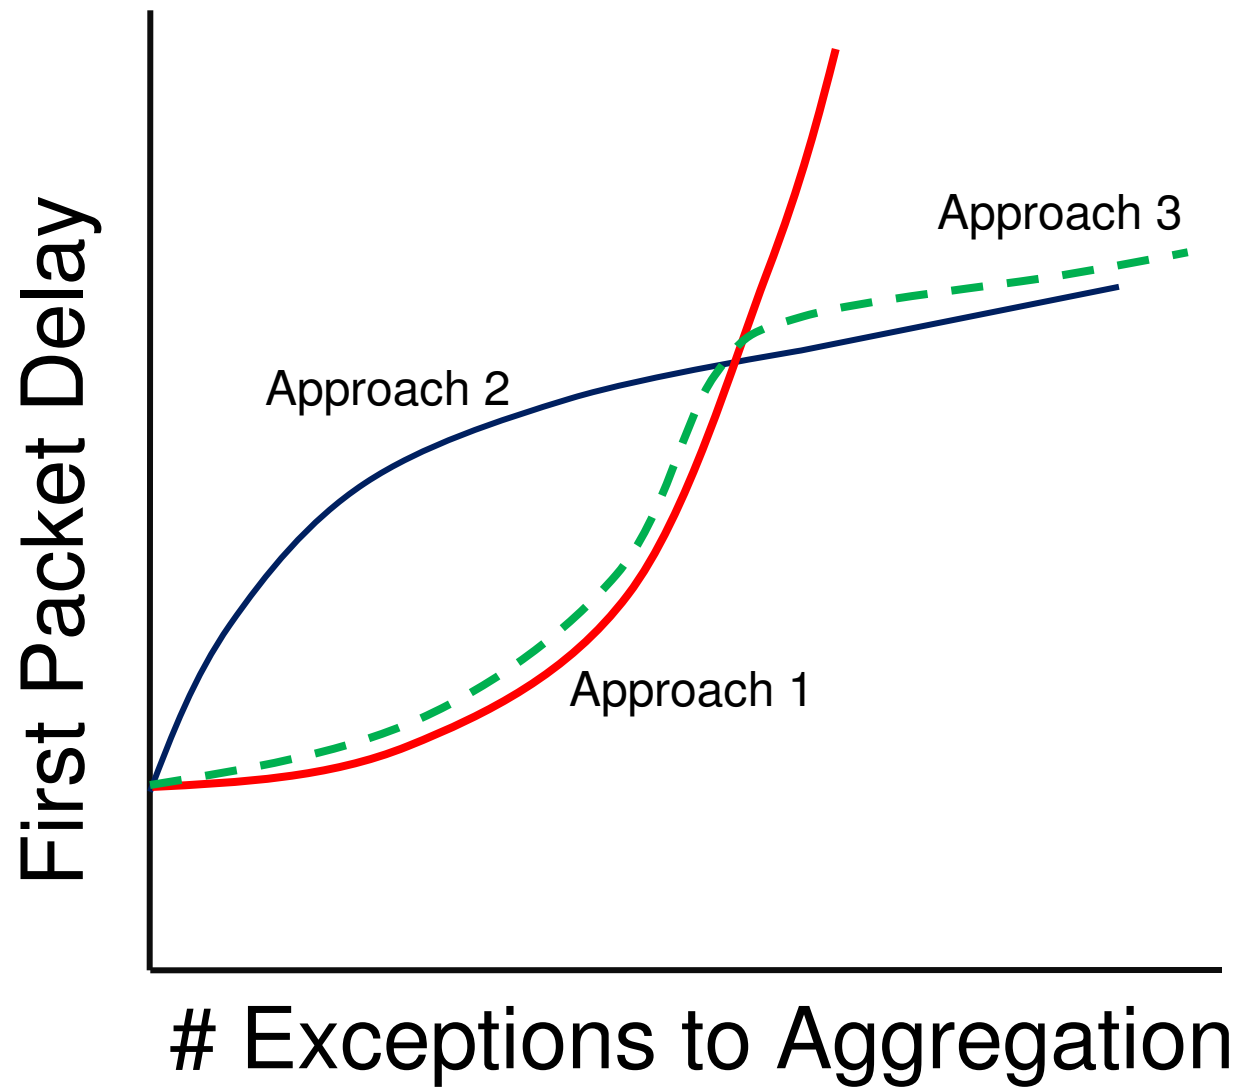
First Packet Delay

Exceptions to Aggregation

Approach 1

Approach 2

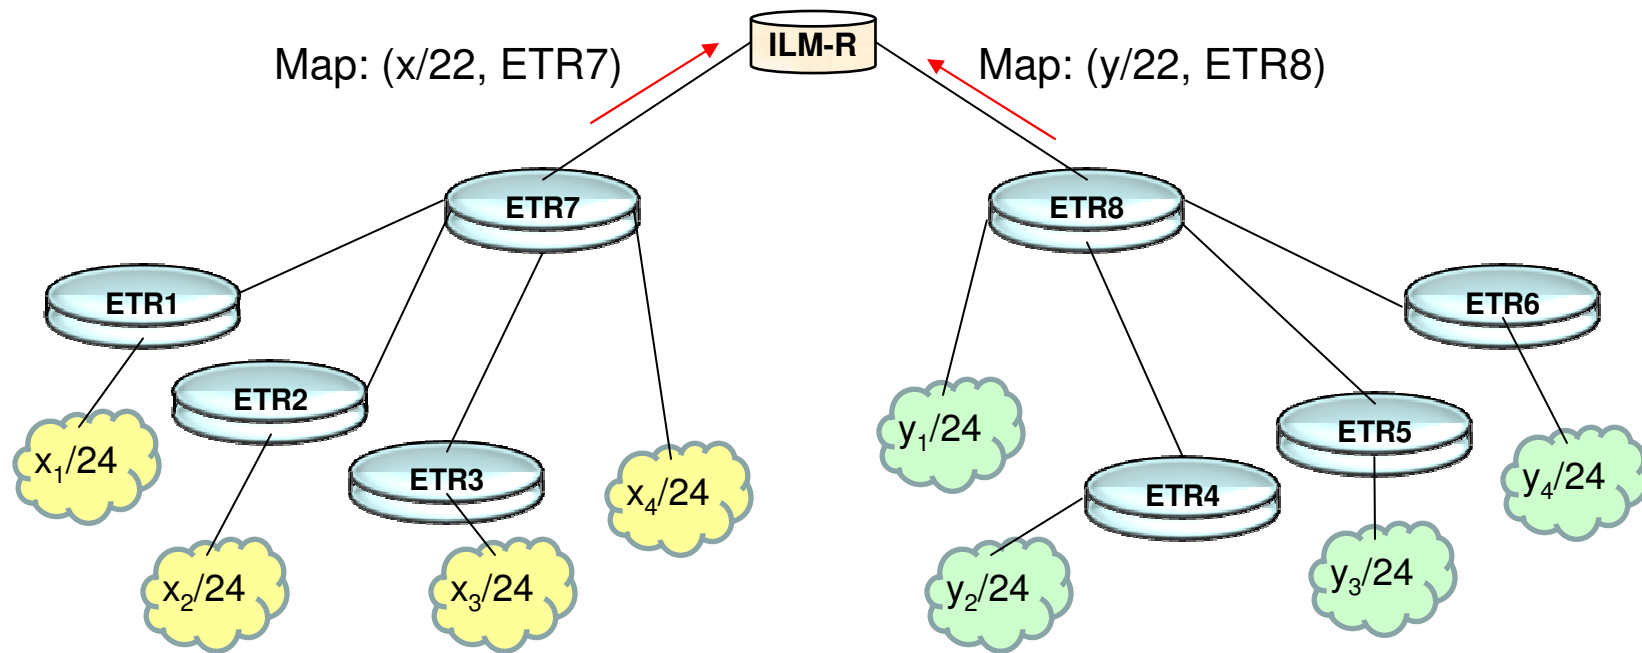Approach 3

Outline of the Talk

- **Problem statement**
- **Management of mapping distribution of subprefixes spread across multiple ETRs**
- **Management of mapping distribution for scenarios with multi-homing and a hierarchy of ETRs**
- **Conclusions / Future Work**

Scenarios with Subnet Multi-Homing

Map: (x/22, ETR7)
Map: (y_1/24, ETR7, backup = yes)

Map: (y/22, ETR8)
Map: (x_4/24, ETR8, backup = yes)

ILM-R

ETR7

Multi-homing

ETR8

x_1/24

x_2/24

x_3/24

x_4/24

y_1/24

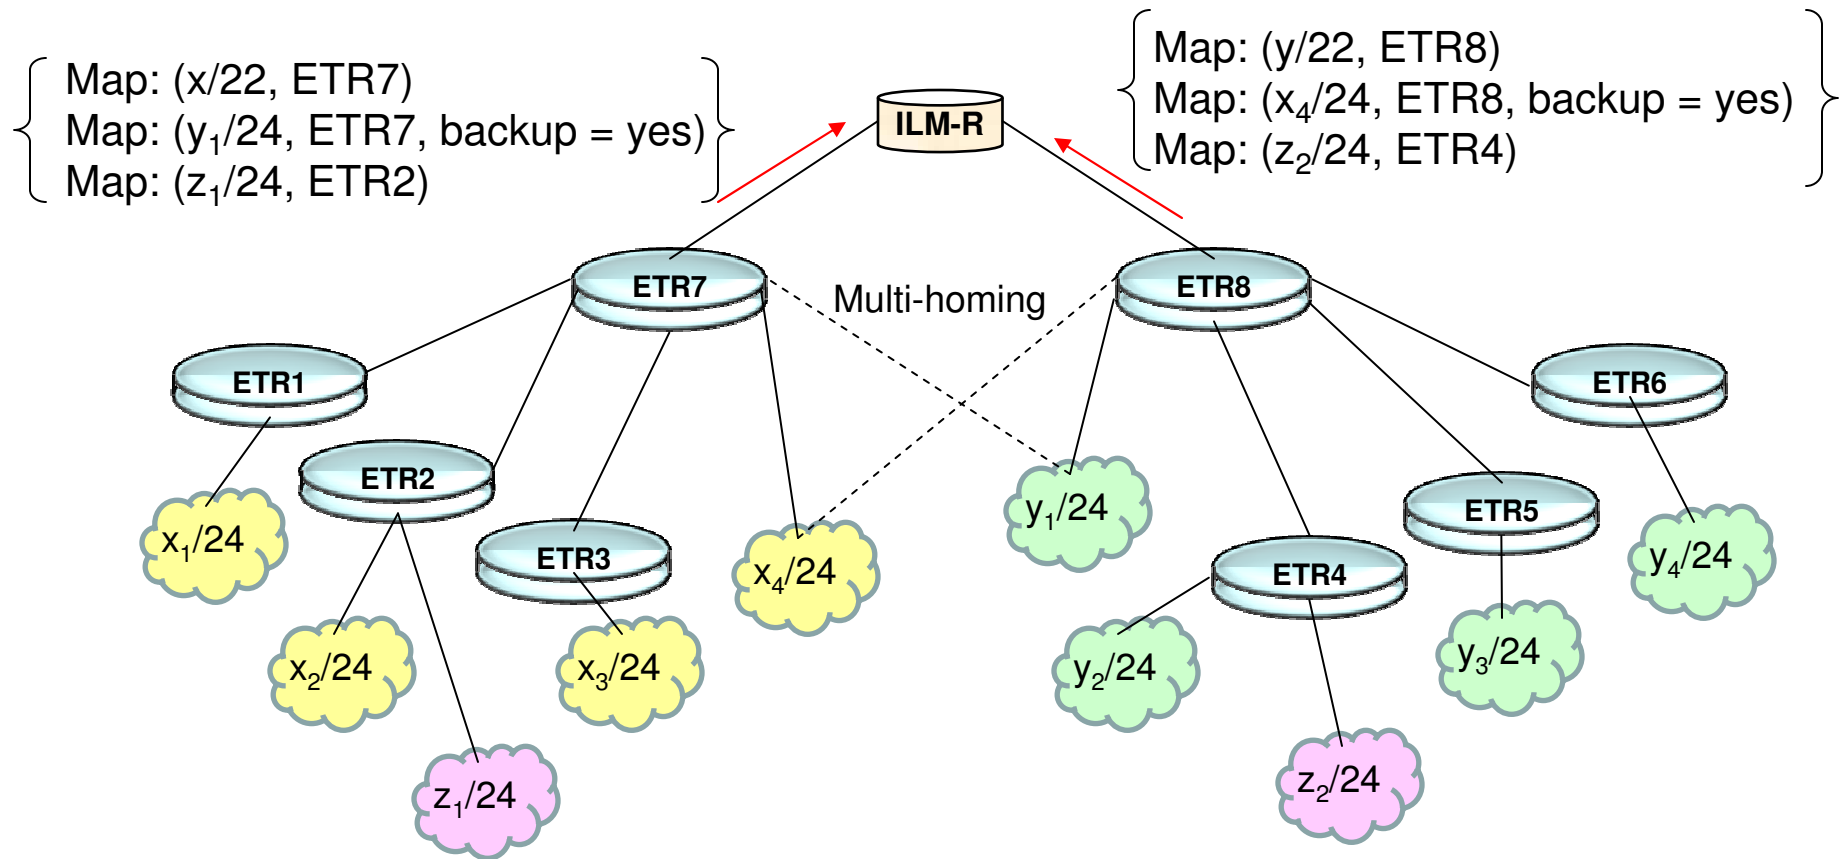y_2/24

y_3/24

y_4/24

- If the backup indication is not provided, then the backup path may take precedence over the primary (longer prefix match)

Scenarios with a Hierarchy of ETRs

Map: (x/22, ETR7)   ILM-R   Map: (y/22, ETR8)

ETR7   ETR8

ETR1   ETR6

ETR2 x_1/24   ETR3 x_4/24 y_1/24   ETR5 y_4/24

ETR4

x_2/24 x_3/24 y_2/24 y_3/24

- Aggregation of EID subprefixes of ETR1, ETR2, ETR3 at ETR7

- Method of ETR7 to { ETR1 / ETR2 / ETR3 } packet forwarding:

  – Table look up: EID address to next hop?

  – Consideration of recursive map-and-encap?

Scenarios with a Hierarchy of ETRs: Combination of Possibilities

Map: (x/22, ETR7)
Map: (y_1/24, ETR7, backup = yes)
Map: (z_1/24, ETR2)

Map: (y/22, ETR8)
Map: (x_4/24, ETR8, backup = yes)
Map: (z_2/24, ETR4)

ILM-R

ETR7

ETR8

Multi-homing

ETR1

ETR6

ETR2

y_1/24

ETR5

x_1/24

ETR3

x_4/24

ETR4

y_4/24

x_2/24

x_3/24

y_2/24

y_3/24

z_1/24

z_2/24

Conclusions and Future Work

- Discussed various architectural questions related to the mapping distribution and management

- Raised some questions regarding aggregation possibilities for the EID address space associated with delivery networks

- Ways of dealing with exception subnets or subprefixes in the mapping distribution protocol

- Consideration of recursive map-and-encap